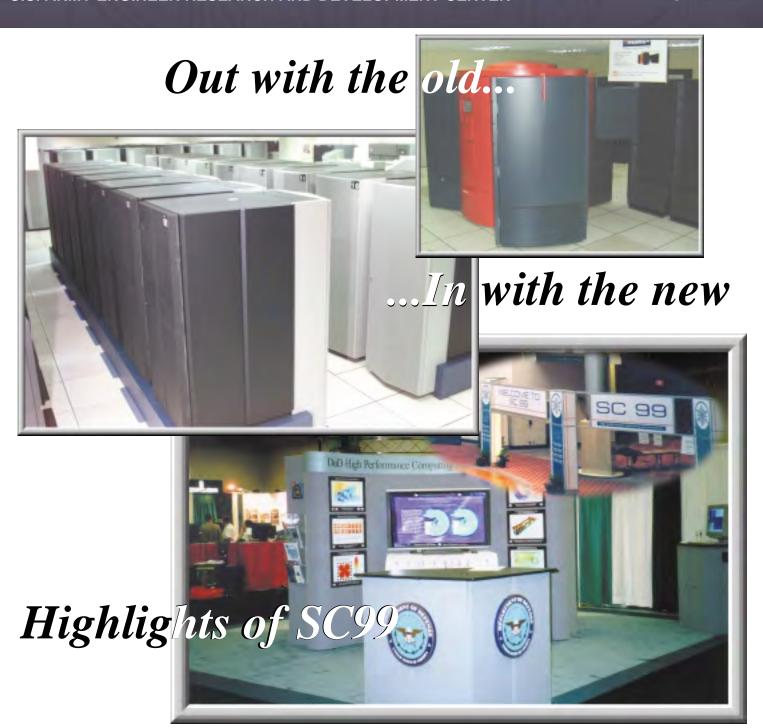
The

e RESOURCE

U.S. ARMY ENGINEER RESEARCH AND DEVELOPMENT CENTER

NEWSLETTER



Major Shared Resource Center







maintaining the data needed, and coincluding suggestions for reducing	ection of information is estimated to ompleting and reviewing the collect this burden, to Washington Headqu ald be aware that notwithstanding an OMB control number.	ion of information. Send comments arters Services, Directorate for Info	s regarding this burden estimate ormation Operations and Reports	or any other aspect of the s, 1215 Jefferson Davis	his collection of information, Highway, Suite 1204, Arlington	
1. REPORT DATE 2000		2. REPORT TYPE		3. DATES COVE 00-00-2000	ERED 0 to 00-00-2000	
4. TITLE AND SUBTITLE			5a. CONTRACT NUMBER			
The Resource. Spri	ng 2000	5b. GRANT NUMBER				
				5c. PROGRAM E	ELEMENT NUMBER	
6. AUTHOR(S)					5d. PROJECT NUMBER	
					5e. TASK NUMBER	
					5f. WORK UNIT NUMBER	
U.S. Army Enginee	rce Center,3909 Ha	velopment Center,A	ATTN: ERDC	8. PERFORMING REPORT NUMB	G ORGANIZATION BER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)					10. SPONSOR/MONITOR'S ACRONYM(S)	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAIL Approved for public	ABILITY STATEMENT	on unlimited				
13. SUPPLEMENTARY NO	TES					
14. ABSTRACT						
15. SUBJECT TERMS						
16. SECURITY CLASSIFIC		17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON		
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	Same as Report (SAR)	32		

Report Documentation Page

Form Approved OMB No. 0704-0188



As we develop this newsletter during the millennium "mania," we naturally reflect on as well as forecast the industry. The U.S. Army Engineer Research and Development Center (ERDC) Major Shared Resource Center (MSRC) recently decommissioned the first supercomputer purchased as part of the High Performance Computing Modernization Program (HPCMP), the Cray C90. At the same time, it is installing one of the largest and newest scalable parallel systems in the world—the IBM Power3, 8-way SMP. The HPCMP has truly modernized the

computing capability for the Department of Defense (DoD) researcher. However, we must continue to stay on the leading edge by forecasting and preparing for the next generation of high-performance computing (HPC) architectures. We must continue to work toward meeting all of our user-specified requirements.

SC99 revealed that clusters of computers are starting to show their strength in the industry; Java may actually become a programming language that will give Fortran and C a run for their money; and long-haul networks are now truly capable of supporting distributed-computing applications. These trends share one common characteristic that foretells the next supercomputing generation—the business of supercomputing is going to be geographically distributed. The biggest challenge associated with distributed-computing environments is "interfacing." We in the computer industry typically think of interfacing as a technical challenge. How often have you heard the term "the standard interface"? If we ever thought technical standards were difficult to establish, imagine interfacing multiple organizations, a fundamental requirement for the success of distributed-computing environments. High levels of data integrity, security requirements, multidiscipline projects, and growing HPC requirements all demand that the supercomputing, researchand-development, and test-and-evaluation communities interface at a level never before experienced. In the spirit of this challenge, we will have the opportunity to reaffirm and strengthen our existing relationships while fostering new ones. These are relationships that will build the "standard interface" of our future.

> Bradley M. Comes Director, ERDC MSRC

About the Cover:

Out with the old...in with the new. While October 1999 marked the end of an era in vector processing at the ERDC MSRC with the removal of the Cray C90, January 2000 marked the beginning of a new era—that of scalable, symmetric multiprocessing. The ERDC MSRC introduced into its capabilities one of IBM's latest and most powerful scalable architectures—the Power3 SMP (see story, page 8).

SC99. In the long-standing tradition of the SCxy supercomputing conferences, SC99 outdid its predecessors. There were more vendors, more floor space, more tutorials, faster computers, and real-time capability for closed captioning and sign language (see story, page 14).

Features

Four Different Parallel Architectures - Which One Is Best?





SC99:

Technologies to Transform the Future

14

ERDC MSRC Contributions to SC99

15

8





SC2000:

The Real Millennium Conference

19

Meet Your Customer Assistance Team

20



Regulars

Announcements	2
Off-Campus	4
Visitors	6
Interview withDr. Joseph Werne	7
Take the Quiz	12
Community Outreach	13
Tool Tips	22
Acronym List	27
In Perspective	28
Technical Reports	28



ERDC MSRC
Computer Systems
Are Gems

25



New ERDC MSRC Web Site

In December, the ERDC MSRC launched a new version of the MSRC's web site (same web site address: www.wes.hpc.mil). We've updated and reorganized our information for easier access, and we have an entirely new navigation system. We'd like to hear your comments on our new look and feel. Please e-mail the webmaster once you've browsed through the site and give us your thoughts and suggestions.

Robert (Jay) Sykes Promoted System Administrator

Jay is our newest system administrator. He has been on the ERDC MSRC team as a member of the Customer Assistance Center (CAC) for several years and has achieved a reputation as an onsite Kerberos and portable batch system (PBS) expert within the ERDC MSRC community. His initial responsibilities in his new position are serving as administrator for onsite UNIX workstations and accounting for the HPC systems.

Nichols Research Corporation Merges with Computer Sciences Corporation

On November 17, 1999, Nichols Research Corporation—the systems integrator for the ERDC MSRC—merged with Computer Sciences Corporation (CSC), an international, \$8 billion company specializing in information technology (IT) services. Their capabilities include E-business strategies, management and IT consulting, systems development and integration, application software, and IT and business process outsourcing. With the merger, CSC now has 57,000 employees.

With Nichols' HPC capability for the science and engineering community, CSC is creating an HPC Center of Excellence to address the outsourcing trend toward more integrator participation across the Federal market. CSC is also the integrator for the Aeronautical Systems Center (ASC) MSRC, located at the Wright-Patterson Air Force Base (WPAFB), Ohio.

Users Group Conference

The DoD HPCMP Users Group Conference 2000 will be held at the Hilton Albuquerque, Albuquerque, NM, June 5-9, 2000.



Paul A. Shahady Retires

Mr. Paul A. Shahady, Director of Information Technology for the ASC, WPAFB, Ohio, retired on December 31, 1999, after 31 years of Federal service. Mr. Shahady was the Chief Information Officer for the WPAFB, Director of the ASC MSRC, and Director of the Simulation and Analysis Facility. He entered the Senior Executive Service in 1997 and was the senior Air Force civilian on the Defense Department's HPCMP team and significantly contributed toward the building of a strong collaborative environment between the Program's elements, including the ERDC MSRC.

While at ASC, Mr. Shahady held a variety of positions on the Air Force Materiel Command headquarters staff, including Director of Reengineering Support and Technical Adviser to the Director of Communications and Information. He also served as Technical Director for Systems with the Logistics Systems Business Center of the Defense Information Systems Agency. He held several positions in the Aeronautical Systems Division, including Director of Information Resources Management and Technical Director of the Information Systems and Technology Center.

Mr. Shahady was commissioned as a Distinguished Graduate of Air Force ROTC from Catholic University of America in 1965 and began his Federal career at the Aero Propulsion Laboratory, WPAFB, in 1968.

Multiphysics Workshop

Modeling of the environment is by nature a multidisciplinary problem requiring the simulation of a number of physical phenomena. The solution of large-scale environmental quality problems requires coupling of physical models and the resources of high-performance computing. Though much progress has been made in recent years, there are still many mathematical and computational issues in coupling different physical models and simulators.

To address the above issues, a workshop entitled "Workshop on Coupling Multiphysics Problems in Environmental Simulation" was held at the ERDC, Vicksburg, MS, on January 11-12, 2000. The workshop was sponsored by the DoD High Performance Computing Modernization Program's ERDC MSRC Programming Environment and Training (PET) program. Organizers of the workshop included Professor Mary F. Wheeler, The University of Texas at Austin, Professor Keith Bedford, Ohio State University, and Dr. Phu Luong, ERDC MSRC Environmental Quality Modeling and Simulation (EQM) onsite PET Lead.



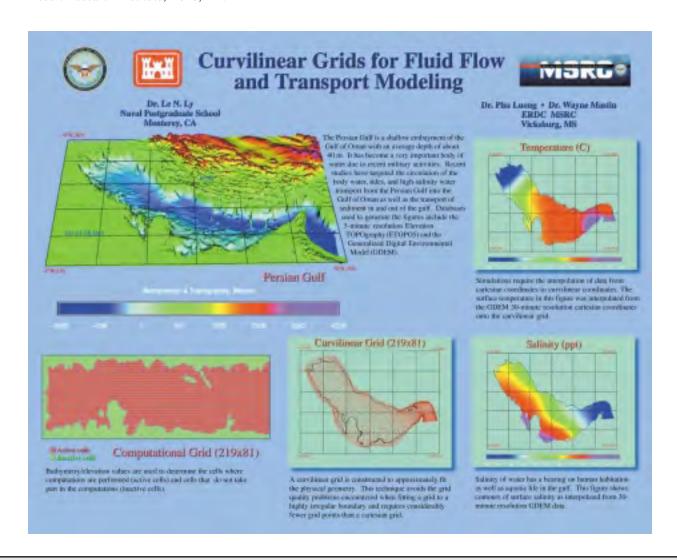
Professor Mary Wheeler, The University of Texas at Austin, addresses a group of researchers at the recent "Multiphysics Workshop" held at the ERDC MSRC. Thirty-one presenters and attendees were present from universities and Government agencies for the 2-day workshop.

PET Team Presents Use of Curvilinear Grids to Model the Persian Gulf and Monterey Bay at the 8th International Meshing Roundtable

Dr. Wayne Mastin, onsite PET Academic Team Lead for the ERDC MSRC, presented a poster entitled "Curvilinear Grids for Fluid Flow and Transport Modeling" at the 8th International Meshing Roundtable in South Lake Tahoe, CA (October 10-13, 1999). Dr. Phu Luong (first author), ERDC MSRC PET Lead for EQM, and Dr. Le N. Ly, Naval Postgraduate School, were coauthors.

The poster describes the application of curvilinear grids for modeling circulation in bodies of water with complicated topography. In this presentation, the Persian Gulf is shown with a computational grid and bathymetric contours. Plots of surface temperature and salinity obtained from available databases are also included. This information can then be used in a simulation program for circulation and transport modeling.

Curvilinear grids were also applied to Monterey Bay, as described in a recent presentation at the Third Conference on Coastal Atmospheric and Oceanographic Prediction Processes held November 3-5, 1999, in New Orleans, LA. The paper presented was "Response of the Monterey Bay Region to Wind Forcing by an Atmospheric Model," Dr. Le Ngoc Ly and Dr. J. D. Paduan, Naval Postgraduate School, Dr. Phu Luong, ERDC MSRC, and Dr. D. Koracin, Desert Research Institute, Reno, NV.



Recent Publication

Ly, L. N., and Luong, P. 1999. "Numerical multi-block grids in coastal ocean circulation modeling," *Applied Mathematical Modeling*, Vol 23(11):865-879.

ERDC MSRC Staff Participates in Visualization Conferences

John E. West

Visualization '99

Dr. Michael M. Stephens, Ms. Christine E. Cuicchi, and Mr. John E. West, ERDC MSRC, attended the 10th annual IEEE (Institute of Electrical & Electronics Engineers) Visualization Conference held October 24-29, 1999, in San Francisco, CA. The conference continued its long tradition of showcasing leading edge research in scientific visualization and data exploration. The evolution of scientific visualization over the past 10 years was evident in the selection of featured speakers. The keynote address featured Paul Smith and John van Rosendale, U.S. Department of Energy, who summarized the results of the series of workshops that led to the creation of the Data and Visualization Corridors Program. They also shared the Program's vision for a broad-based research agenda aimed primarily at facilitating the manipulation and exploration of very large-scale data sets. The capstone address by Rick Stevens, Argonne National Laboratory and The University of Chicago, focused on the ActiveSpaces project, which creates integrated wholeroom visual environments for exploring data. These rooms combine a range of display, tracking, and communication systems linked to the National Grid to create environments for large-scale collaboration among groups of users. The keynote and capstone addresses reflect the current thinking in the visualization community that the path through the challenges created by very large data sets is through complex, integrated systems—a theme found repeatedly throughout the conference.

John West was also the coauthor of a technical paper presented during the conference that described a new technique for objective comparison of data and images. The paper, entitled "Structured Spatial Domain Image and Data Comparison Metrics," was the result of collaborative research between the ERDC MSRC and the Mississippi State University (MSU)/ National Science Foundation (NSF) Engineering Research Center. Contributing to the research were Nivedita Sahasrabudhe, MSU, John West, ERDC MSRC, Professor Raghu Machiraju, MSU, and Professor Mark Janus, MSU.

InfoVis and PVG

The visualization conference was preceded by two focused events: the Symposium on Information Visualization (InfoVis) and the Symposium on Parallel Visualization and Graphics (PVG). In its fifth year, InfoVis '99 continued its unique focus on techniques for understanding data with no underlying physical representation, a class of data that continues to grow dramatically along with the tremendous growth of computer applications for medicine, instrumentation, and e-commerce. The inaugural year for PVG focused on the parallelization-of-graphics techniques on commodity PC clusters.

The ERDC MSRC offers visualization capabilities for high-performance computing. If your current visualization processes are becoming overwhelmed by your computational modeling techniques, make sure to contact the MSRC.



Christine Cuicchi (left), John West (center), and Dr. Michael Stephens (right) attended the 10th annual IEEE Visualization Conference.

Robert W. Numrich

On October 13, 1999, Dr. Robert W. Numrich, Principal Scientist at Cray Research, Inc., conducted a half-day seminar on Co-Array Fortran. Co-Array Fortran is a parallel extension to Fortran 90/95. It converts Fortran into a robust, efficient parallel language and requires Fortran programmers to learn only a few new rules. These new rules are related to two fundamental issues that any parallel programming model must resolve: work distribution and data distribution. Co-Array Fortran uses a simple, explicit notation for data decomposition, such as is often used in message-passing models, expressed in a natural Fortran-like syntax. The syntax is architecture-independent and may be implemented on distributed-memory machines, shared-memory machines, and even on clustered machines. For some applications, it can simplify parallel programming. Co-Array Fortran is currently available on the ERDC MSRC Cray T3E.

The first part of the seminar consisted of a lecture on Co-Array Fortran syntax, followed by coding examples. Afterwards, Dr. Numrich was available for individual consultation. Additional information, documentation, reports, and example code can be found on the Co-Array Fortran web site: www.co-array.org.



Dr. Robert Numrich

David Stinson discusses HPC with

Leadership Vicksburg.

U.S. Government Impact on Local Community

The Vicksburg-Warren County Chamber of Commerce, Vicksburg, MS, created Leadership Vicksburg, a group of annually chosen key individuals who will be most involved in shaping Vicksburg-Warren County's future, to bring together participants with diverse backgrounds and experiences who have demonstrated talent and commitment for leadership and to give them direction and support.

The Leadership Vicksburg session for this year entitled Federal Impact

included a tour of the ERDC, Vicksburg, the U.S. Army

Engineer District, Vicksburg, and the National Military Park. Leadership Vicksburg participants became more aware of the Government impact on the economy and social/political structure of the Vicksburg community.

On January 22, 2000, Mr. David Stinson, PET Monitor and Outreach Program Manager of the ERDC MSRC, and Ms. Christine Cuicchi, computational engineer for the ERDC MSRC, conducted a tour of the HPC facility and the Scientific Visualization Center as part of the Leadership Vicksburg's visit to the ERDC Information Technology Laboratory (ITL).



Christine Cuicchi demonstrates scientific visualization using the ImmersaDesk.

Chief Information Officer and Deputy Chief of Staff for Corporate Information, U.S. Army Corps of Engineers, Visits ERDC, Vicksburg

The new Chief Information Officer and Deputy Chief of Staff for Corporate Information, U.S. Army Corps of Engineers, Mr. Wilbert Berrios, visited ERDC, Vicksburg, MS, on November 16 and 17, 1999.

COL Robin R. Cababa, Commander, ERDC, welcomed Mr. Berrios and presented an overview of ERDC. Mr. Tim Ables, Acting Director, ITL, served as Mr. Berrios' escort for his visit to ERDC, Vicksburg.



COL Cababa, Commander, ERDC, (left) welcomes Wilbert Berrios (seated) to ERDC, Vicksburg.

(Continued next page)

While touring ITL, Mr. Berrios visited the ERDC MSRC computing facility and Scientific Visualization Center, where Dr. Carl Cerco, Environmental Laboratory, ERDC, demonstrated one of his projects—Chesapeake Bay Model—at the ImmersaDesk, and Dr. David Horner, Geotechnical Laboratory (GL), ERDC, demonstrated his former Challenge Project—Grizzly Mine Plow—on the Panoram system in the Collaboratorium.



Wilbert Berrios (left) with Tim Ables, Acting Director, ITL, in the ERDC MSRC's computing facility

interview with . . .

Dr. Joseph Werne has been a DoD Challenge Project user at the ERDC MSRC for the last 3 years. His work at Colorado Research Associates has brought substantial contributions to the Air Force AirBorne Laser program and to the field of computational fluid dynamics. We spoke with him about his research and the computational challenges involved. His current allocation on the ERDC MSRC HPC systems is 350.000 node-hours.

What are you modeling? Dave Fritts and I at Colorado Research Associates are simulating turbulence in the lower stratosphere and upper troposphere in support of the AirBorne Laser (ABL) program. The ABL project tackles a monumentally challenging design problem: to focus a coherent, high-power laser beam on a missile target from up to 200 km away. The major design hurdle to overcome is the distortion and defocusing the laser beam experiences as it propagates through turbulence in the atmosphere. The problem is similar to that encountered when making astronomical observations with earth-based telescopes, which must look up through the turbulent atmosphere. However, ABL design is much more challenging because it must be capable of operating along much longer paths through the atmosphere, which may contain significant levels of turbulence. ABL will deal with the distorting effects of turbulence in the atmosphere by intentionally defocusing the laser beam to compensate for the turbulence it encounters. The idea is to treat turbulence in the atmosphere as a complicated lens that refocuses the intentionally defocused beam. The strategy with which one compensates the laser beam for atmospheric turbulence based on what one learns by viewing the distorted target is known as "adaptive optics." The challenge for ABL design is to anticipate the spatial distribution and intensity of turbulence expected in the atmosphere and to deduce its specific impacts on laser-beam propagation. That's the focus of our work.

Dr. Joseph Werne ERDC Challenge User

The problems in turbulence theory and application are challenging, and after 2 years we are only just now positioned so that ABL design may fully benefit from the simulations we have conducted. There is a clear need to continue our ABL Challenge effort.

What are the computational challenges? Atmospheric turbulence spans an enormous range of length and time scales. If one were to attempt to simulate it directly, without any approximation to lessen the computational burden, spatial resolution on the order of 20,000 x 20,000 x 20,000 grid points would be required, and several billion central processing unit (CPU) hours on the ERDC Cray T3E would be needed to finish just one simulation—for a region of the atmosphere spanning just 1 km in each spatial direction.

Two standard approaches are employed to combat this computational challenge: (a) conduct so-called largeeddy simulations that incorporate subgrid-scale (SGS) models intended to mimic the effects of small-scale turbulence not resolved by the numerical solutions, or (b) employ the so-called scale similarity exhibited by atmospheric turbulence to extrapolate accurate solutions of less-turbulent simulations to the levels of turbulence attained in the atmosphere. Method (a) is used in many engineering applications. Unfortunately, we cannot employ it here because reliable SGS models have not been developed that properly account for the effects of density layering (or stratification) present in the atmosphere. As a result we exploit method (b), which can only be properly implemented by conducting the highest resolution and largest possible simulations on the ERDC MSRC's Cray T3E. The method requires not just one simulation but a succession at finer and finer resolution, which attain higher and higher turbulence intensities. The simulations we have conducted at the ERDC and NAVO MSRCs represent

(Continued page 26)



Four Different

SGI Origin

Cray T3E



IBM SP



IBM Power3 SMP

A frequently asked question at the ERDC MSRC is, "Which HPC system is the best?" A better question is, "Which system is best for my application?"

To answer the latter question, one must be aware of the strengths, weaknesses, and differences among each of the four parallel architectures at the ERDC MSRC. This section first looks at some HPC history and then examines each of the four ERDC MSRC HPC architectures.

Parallel Architectures

- Which One Is Best?



2000

Perspectives of HPC

Since the advent of the von Neumann architecture model, computer architectures have not changed significantly over the last 50 years. The major components to be found in such computers are some form of Arithmetic-Logic Unit (ALU) to perform the computations' memory storage for holding data currently being processed

and input/output (I/O) devices to communicate with the outside world. As technology has improved, processing speeds have become faster; memory sizes have increased and often use hierarchical designs; and I/O devices have become smaller in size, larger in capacity, and cheaper to build.

In the quest to process more data more quickly, vector and parallel have been two commercially successful classes of computer architectures. Vector processors take advantage of the fact that even simple arithmetic operations performed on a computer are broken down into multiple stages to compute the result. In a sequentially executing computer, the start of a computation must wait for the completion of all stages of the previous one. However, if a computation repeatedly executes the same operation—for example, addition of corresponding elements of two vectors—special ALU hardware can be constructed such that after the first pair of values has completed the first stage of the addition process, the second pair of values may begin the first stage while the first pair proceeds to the second stage. The third pair of values follows closely after the second pair, and so on. Thus, results are produced in the time it takes to execute one stage of the calculation rather than having to wait for all stages to be completed before the next computation can begin.

The compilers for these machines were designed to recognize computations that could be vectorized and make use of the special hardware. Very little code modification was necessary to take advantage of vector processing.

While vector machines could achieve speedups for performing the same operation on multiple operands, the recognition that some algorithms contained code segments that were different but independent was acknowledged. They could be executed at the same time. To take advantage of this type of parallelism,

Dr. Clay P. Breshears

multiple processors could be employed. For instance, when adding two vectors, the result from the addition of the first pair of elements has no effect on the addition of the tenth pair of elements or any other pair. Thus, different sequential processors can be used to perform the addition of different portions of the two vectors. Because no special modifications need to be made to the processors that constitute a parallel computer, these architectures are easier to design and may be built from less expensive, commercially produced components.

The first multiprocessing computers were built with a shared-memory architecture. That is, each processor had access to all the memory, typically through a shared bus. Shared-memory multiprocessors are easily programmed because each processor is able to read or write to any variable in the program. Unfortunately, the bottleneck for these systems was the shared bus used to access the global memory space. This ultimately became a significant design constraint that limited the practical size of shared-memory parallel computers to several dozen processors.

Distributed-memory architectures worked around this limitation by attaching a small amount of memory to each processor. To access data stored in the memory of another processor, a message had to be sent from the processor holding the data to the processor needing the data. While distributed-memory machines were no longer restricted to small numbers of processors, there were restrictions based on the processor-to-processor network configuration, such as the number of connections that could be supported by each processor, the distance (number of intervening processors) between any two processors in the system, and the capacity of the network wires. Different network topologies gave different performance characteristics, and different types of applications performed better on different topologies. Also, distributed-memory machines were much more difficult to program because every instance of data sharing required the programmer to explicitly code the interprocessor communication.

The new parallel computers being designed and built today are a combination of the best parts of shared-and distributed-memory architectures. A handful of processors are clustered together around a shared memory. These clusters are then connected to each other via a network.

ERDC MSRC Hardware

The ERDC MSRC currently has four different parallel computers. With the removal of the Cray C90 in October 1999, no vector processors remain at the site. The principal differences between the ERDC MSRC parallel platforms are the speed of their processors, their memory size, and the way that their processors interact during a parallel computation. These differences can dramatically affect how well an application runs on a given platform.

Let's examine each of the ERDC MSRC HPC machines with an eye toward how the processors share data and the types of communication patterns that can be found in parallel codes that perform best on each platform.

Cray T3E

The Cray T3E is a tightly coupled, distributed-memory architecture. Each processor in the machine is connected via a three-dimensional (3-D) torus mesh. Processors in a 2-D mesh are directly connected to four other processors (usually designated as North, East, South, and West). A 3-D mesh stacks 2-D meshes on top of each other and connects corresponding processors between each layer adding two more direct connections (Up and Down). A torus wraps the connections of the processors on the edges of the mesh back to the processors on the opposite side of the mesh. For example, the North neighbor of a processor on the North edge of the mesh is the corresponding processor on the South edge and vice versa.

Such a network is good for regular communication patterns, especially if all communication can be performed between processors that are directly

out with the old...

In a move away from the traditional workhorse vector architecture, the ERDC MSRC retired its last vector machine, the Cray C916 (C90), in October 1999. It served for the last 8 years. It was the first supercomputer purchased as part of the HPCMP. In preparation for the retirement, the ERDC MSRC formed the Computational Migration Group



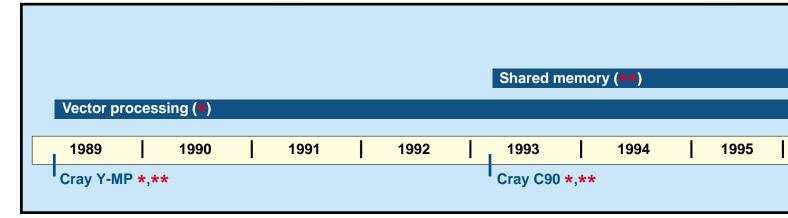
(CMG) to help users migrate their traditional vector codes from the C90 to the other scalable parallel machines in the Center (please see "A New Addition to the Family: The IBM Power3 SMP" in *The Resource*, Winter 1999). Using an internally developed software taxonomy, along with monthly utilization data, the CMG migrated codes from the vector to scalable environments beginning with codes most heavily used on the C90 and moving through the list until all significant codes had been migrated. To date, the ERDC MSRC is unaware of any loss in user productivity because of the removal of its last vector system.

connected to one another. Very few application codes can sustain such restrictive communication. Given a regular communication pattern in a code, though, the communication time should remain constant on the T3E, regardless of problem size. Global communications (e.g., broadcast, reduce, gather/scatter) can also be optimized to make efficient use of the regularity of the network topology.

IBM SP

The IBM SP is a loosely coupled, distributed-memory system. It was envisioned to be a collection of independent workstations connected through a high-speed network switch. The system could be used as many different, individual computers, or utilizing a message-passing library, some or all of the processors could be harnessed to execute a single, parallel computation.

A Brief History of HPC at the ERDC MSRC



...in with the new

While October 1999 marked the end of an era in vector processing at the ERDC MSRC with the removal of the Cray C916 (C90), January 2000 marked the beginning of a new era—that of scalable, symmetric multiprocessing (SMP). The ERDC MSRC introduced into its capabilities one of IBM's latest and most powerful scalable architectures—the Power3 SMP.

- Mike McCraney



In January 2000, the ERDC MSRC installed a 64 high-node Power3 SMP system. This 512-processor system will complement the existing IBM SP systems, the SGI Origin 2000, and the Cray T3E and will increase the MSRC's overall peak computing capability to 1.4 Tflops. Each of the 64 Power3 high nodes are configured with eight 222-MHz RS6000 Power3 processors, 4 GBytes of memory, and 9 GBytes of local disk.

The nodes are initially connected with a 4-way, 175-MByte/sec interconnect, which will be upgraded to an 8-way switch in mid-2000. The SMP follows the historically successful IBM SP line of scalable parallel systems and is anticipated to be one of the MSRC's workhorses of the future. The system is available to DoD Challenge users as well as to the general-user community. Users interested in learning more about this new system or obtaining an account on any of the systems at the ERDC MSRC should contact the CAC (1-800-500-4722, 7:00 a.m. - 7:00 p.m. Central Daylight Time, Monday through Friday, or by e-mail [info-hpc@wes.hpc.mil]), or visit the ERDC MSRC web site at www.wes.hpc.mil.

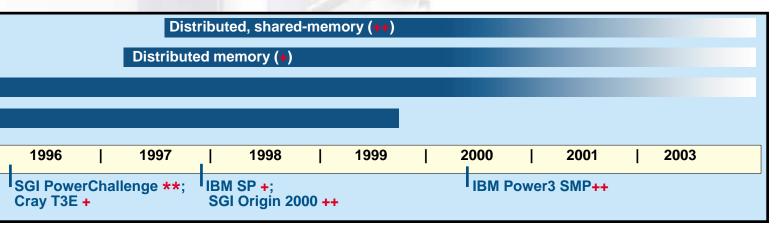
Messages are always routed through the switch network. The physical proximity of processors has little bearing on communication speed. While the time required to send a message between processors that physically reside in the same cabinet may seem longer than necessary, messages between processors that are not located in close proximity to each other should take roughly the same amount of time. The switch is designed to offer alternate routes for messages when there is a heavier amount of traffic. Thus, irregular message-passing patterns are not adversely affected by network topology or number of messages being passed.

SGI Origin 2000

The SGI Origin 2000 is a distributed shared-memory system. The hardware is designed like a distributed-memory architecture where processors are connected through a network and have a separate and private amount of memory. However, the system keeps track of

which memory space holds which variables. So, parallel programs can be developed on the SGI Origin 2000 using either a shared- or distributed-memory model. The programmer may use message-passing interfaces (MPIs) for distributed-memory programming, OpenMP or threads for shared-memory programming, or some combination of both to best suit the application.

Programmers must be aware that the initial distribution of data among the processors for shared-memory computations is important. All data used in a program will tend to be placed within a single processor's memory. In this case, each processor must compete with all the others for access to the shared data, which quickly becomes a bottleneck for the computation. There are commands and programming directives that can be used to distribute the data across the processors. There is also the "first touch" rule, which sets the condition that the first processor to access a data segment will hold that data throughout the execution.



Initializing an array in parallel (using threads, for example) effectively distributes the data and avoids the potential bottleneck.

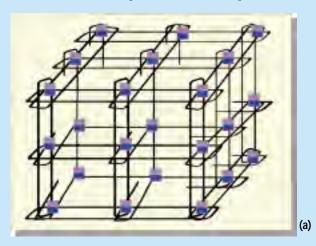
IBM Power3 SMP

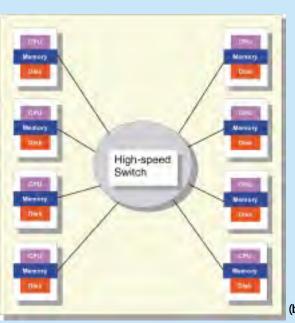
The new IBM Power3 SMP system contains clusters of eight CPU SMP nodes connected by a high-speed network switch. While the SMP nodes have a shared memory, they may also be used as separate message-passing processors within a node or to processors in other nodes. Shared-memory programming works only among processors within the same SMP node.

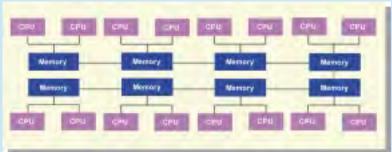
If an application code and data will fit within the memory of a single SMP node, all message passing or data sharing will be confined to the processors and memory of that node. This will effectively ensure that no other processes on the machine will interfere with the execution of that code as they may on the other systems. Thus, even if more processors or memory resources are needed than are available on a single node, codes that can keep message traffic between nodes to a minimum will run effectively.

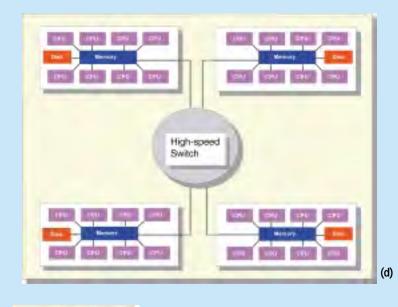
Different machines perform differently on different codes, so choose the best system for given computations.

Question: Can you match the architecture diagrams with the correct parallel computer?









SGI Origin 2000 IBM SP Cray T3E IBM Power3 SMP

LEGEND

- CPU

MEMORY

Answers on pages 22-23

(c)

Introducing HPC to Local High School Students

Mary Gabb

For the past 2 years, Mr. Paul Adams has been visiting a local high school in Clinton, MS, to discuss the HPC field with junior and senior math students. This program began at his suggestion to a friend who teaches at the high school. He travels to the school once every spring and sees approximately 125 students.

Paul's visits generate significant enthusiasm from the students. He describes some of the research projects performed by DoD scientists on HPC systems. In particular, he focuses on Challenge projects because they seem to generate the most interest from the students. Paul also discusses the work involved in maintaining these systems. "I try to give them an idea of how big and how fast our systems are. To make them appreciate what a gigaflop actually is, I tell them that our systems are about 1,000 times faster than the PCs sitting on their desks. It's a rough estimate, but they get the idea."

According to Paul, the students' questions are eclectic, ranging from career development ("What's the outlook for the information technology field?" and "How much do you earn?") to more science-oriented questions. They want to know about the operating systems used at a place like the ERDC MSRC and how the math they are learning in high school applies to the real world.



One of the most difficult aspects of discussing HPC careers with high school students is that they want to know how their homework assignments apply to the real world. Paul Adams tries to give specific examples, but it isn't always easy: "It's tough when they ask me about the relevance of imaginary numbers."

Paul described a typical session: "I sometimes can't even get through my whole presentation. I'm 2 minutes into it, and they already have 10 questions! When they ask about real-world applications to their homework, I try to give them specific examples."

What prompted Paul to start this service? Paul recalls his last 2 years in high school, remembering the frustration of underchallenging and seemingly irrelevant courses. His conversations with today's high school students may quell that restlessness and give them a clearer goal on which to focus.

Paul Adams joined the ERDC MSRC in 1994 as an applications analyst in computational fluid dynamics. He has also worked with CAC at various times. He currently serves as Acting Systems and Networks Manager. He has a B.S. and M.S. in aerospace engineering.

ERDC MSRC Staff Member Responds to Emergency Request

Rose Dykes

In addition to its regular trips to ERDC, Vicksburg, the Mississippi Blood Services acted upon a request by employees of the ERDC Structures Laboratory to send

its mobile unit to obtain blood donations for a young woman in the Vicksburg community who was in desperate need of blood in January and February 2000 before having a bone marrow transplant. At least 20 volunteers were needed to demand a special visit of the mobile blood unit.

Mr. David Stinson was among more than 100 blood donors to volunteer. Those who know David were not surprised that he was at the head of the line to donate in this lifesaving effort. He is just that kind of a person. In addition to his full time duties as the PET Monitor and Outreach Program Manager at the ERDC MSRC, David finds time to spend 2 hr each Wednesday night to work with 120 young people in the AWANA

(Approved Workers Are Not Ashamed) program. David is the Director of AWANA, which includes children from 3 years old through the eighth grade.



David Stinson giving blood



SC99: Technologies to Transform the Future

"If you think that when you've got the Earth wired with the Internet you're done—forget it. You've got a long way to go."

Mary Gabb

Lessons on management are not what one would expect to hear in the keynote address at a supercomputing conference. But during Ms. Donna Shirley's 35 years in the aerospace industry, she has acquired numerous "pearls of wisdom" from working in a highly charged, highly technical environment to push back the boundaries of science and technology—all within a tight budget.

Donna Shirley is perhaps most famous for her role as manager of the Mars Exploration Project at the Jet Propulsion Laboratory (JPL). She recounted the development of the Mars Pathfinder Project as a way to illustrate her model of managing technical teams. While she did reference the computational capabilities and achievements during the Mars Exploration Project, it was the method of creating a highly diverse team that captured the audience as she described her philosophy of "managing creativity"—the leadership and focusing of creative teams.

Donna's First Law:
All creative enterprises
are the result of collective
creativity.

People, not groups, are often described as creative. Yet, Ms. Shirley asserts that all creative enterprises are the result of collective creativity (ref: Donna's First Law),

taking full advantage of the creativity of the individual team members. Even geniuses do not operate in a vacuum.

Ms. Shirley reviewed her system for creative enterprises, which at first glance looks like a project life cycle. It consists of nine processes that occur in order and continue processing throughout a project (e.g., build team, generate concepts, achieve alignment). At the center of this cycle—and the force that holds collective creativity together—is communication.

In the long-standing tradition of the SCxy conferences, SC99 outdid its predecessors. There were more vendors, more floor space, more tutorials, faster computers, and real-time capability for closed captioning and sign language. This section of the newsletter reviews the keynote address, the ERDC MSRC participation in SC99 programs, and SC2000 features.

Ms. Shirley referred to communication failure when she discussed the Mars Orbiter problems. She also discussed numerous specific examples of how management is a balancing act (ref: Donna's Second Law):

- ➤ Too little development time, you risk a substandard product. Too much development time, competitors will overtake you.
- ➤ Too little management control, anarchy results. Too much management—well, many of us can fill in the rest.
- ➤ Too little evaluation, the project can lose focus. Too much evaluation, the team can become demotivated.

All of these management principles are applicable to the HPC industry—a highly charged, highly technical industry operating on increasingly tight budgets. Functioning at what

she calls the "edge of chaos," the boundary between the stable, traditional groups in an organization and the dynamic, creative, and "probably rebellious"

Donna's Second Law: Management is a balancing act.

groups—requires balance. Successfully operating at the edge of chaos is essential to the success of the industry.

Ms. Shirley concluded her talk by mentioning a proposed Internet station on Mars. Therefore, when the final communication satellite on Mars is ready, the Internet station will already be there. Always looking to push back the frontiers of science, she discussed the possibility of wiring the entire solar system. Ms. Shirley challenged the audience, "If you think that when you've got the Earth wired with the Internet you're done—forget it. You've got a long way to go." Who's up to the challenge of managing that creative team?

Since retiring from her position as manager of the Mars Exploration Project at the JPL, Donna Shirley has exercised her expertise in management issues garnered over the last 25 years in managing creative enterprises. Specifically, she has written a book on "managing creativity" as a way to foster creativity for developing new technologies and applying those technologies to new challenges. Nevertheless, isn't "managing creativity" an oxymoron? According to Donna, not at all. (Donna currently serves as Assistant Dean of Engineering for Advanced Program Development at the University of Oklahoma.)

ERDC MSRC Contributions to SC99

David Stinson

Each year, the HPCMP participates in SCxy, a large conference where computer hardware and software vendors show off their wares, and universities and research organizations present projects that they have in the works. Every year one of the MSRCs is tasked with organizing a display booth to highlight DoD research projects and the role of HPC within the DoD. Although, only one MSRC is responsible each year for organizing the DoD booth, the other MSRCs and Distributed Centers provide input in the way of graphics, demonstrations, and booth workers. This year, the organization responsible for the DoD HPCMP booth at SC99 in Portland, OR, was the Naval Oceanographic Office (NAVO) MSRC. David Stinson coordi-

nated contributions from the ERDC MSRC to the booth. This involved collecting and formatting video clips, posters, and supporting documentation of the research being conducted on ERDC MSRC HPC resources. These were provided to the NAVO MSRC for inclusion on a compact disk (CD), which was given away with a CD wallet to visitors at the DoD booth. Some of the posters submitted for the CD were incorporated into a PowerPoint display, which was shown on a large flat-panel plasma display at the DoD booth. One of the ERDC MSRC research project submissions selected for display was on work conducted by an ERDC Coastal and Hydraulics Laboratory researcher, Dr. Zeki Demirbilek. Dr. Demirbilek conducted a study of the

wave field around a proposed floating airport to be built in Osaka Bay, Japan. The predicted wave estimates will be used for design calculations and loading, motion, and mooring analyses of the floating airport. Another of the ERDC MSRC submissions was on work conducted by ERDC researchers Dr. David Horner, GL, and Alex Carrillo, MSRC. Dr. Horner's work on the Grizzly Mine Plow and on the development of virtual proving grounds to evaluate vehicle-soil interactions for a variety of military vehicles was highlighted.



David Stinson presents a framed miniature version of an SC99 poster to Dr. David Horner in appreciation for his support of the DoD booth at SC99. Other recipients not pictured include Dr. Zeki Demirbilek and Alex Carrillo.



Mine Plowing Simulation



Wave Climate Around a Proposed Floating Airport in Osaka Bay, Japan

The ERDC MSRC

Mary Gabb

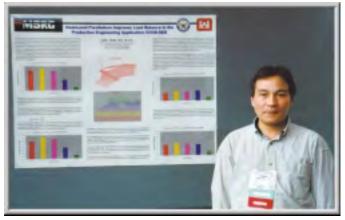
Poster Sessions

Drs. Phu Luong, ERDC MSRC PET EQM Lead, Clay P. Breshears, ERDC MSRC PET Scalable Parallel Programming (SPP) Tools Lead, and Henry Gabb, Director of Scientific Computing, ERDC MSRC, presented a poster at SC99 entitled "Dual-Level Parallelism Improves Load Balance in the Production Engineering Application CH3D-SED." The poster shows how dynamic threading was used with OpenMP to load balance an MPI application. A wide range of scientific topics were covered at the SC99 poster session, including specific applications and four presentations on Beowulf cluster performance. Presenters were given two morning receptions to discuss their work.

Expert Panel Discussions

A panel of six experts met to discuss their experiences with dual-level, or hybrid, parallelism in a panel discussion entitled "Experiences with Combining OpenMP and MPI." Among the panel members was Dr. Henry Gabb. Each panelist spoke for approximately 10-15 min, describing a specific example or examples of how they combined message-passing and parallel compiler directives. Dr. Gabb described work done at the ERDC MSRC on three parallel applications:

➤ A coastal hydrodynamics application. This work was featured in the SC98 Challenge Demonstration. (See *The Resource*, Fall 1998.)



Dr. Phu Luong at poster session. At the SC99 conference, the acceptance rate for posters was only one in four.

- A structural acoustics application that simulates sound waves interacting with a solid structure.
- ➤ A sediment-transport application. This work was also presented as an SC99 poster from the ERDC MSRC (see previous article).

Dr. Gabb described the constraints of working with legacy codes and applications that are used in engineering design projects. For example, the code must be parallelized with as little source code modification as possible. When trying to implement OpenMP and, especially MPI, working around this restriction can be difficult.



The panel members are shown, left to right: Dr. Danesh Tafti, Senior Research Scientist and Associate Director, Application Technologies Division, National Center for Supercomputing Applications; Greg Gaertner, Principal Software Engineer, Compaq Computer Corporation; Dr. Howard Scott, Physicist, Lawrence Livermore National Laboratory; John Levesque, Director of the Advanced Computing Technology Center, IBM Corporation; Dr. Henry Gabb, Director of Scientific Computing, ERDC MSRC; Dr. Stef Salvini, HPC Group Leader, Numerical Algorithms Group, Inc.

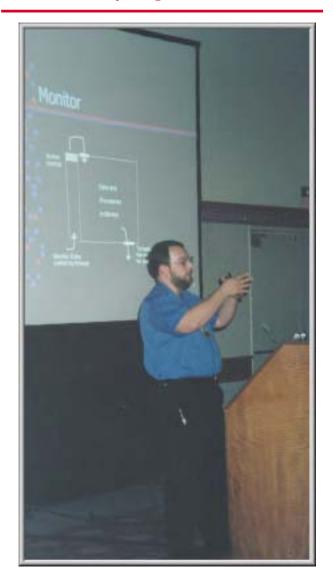
Team at SC99

Expert Panel Discussions (Continued)

According to the panelists, adding OpenMP to an existing MPI application is not that much more difficult than adding OpenMP to a serial program. In many, but not all, of the case studies presented, OpenMP and MPI combined scales better than either model by itself. As Dr. Gabb explained, "Combining MPI and OpenMP allows us to throw more processors at the problem." Two case studies showed how dynamic multithreading can be used to improve the load balance in an MPI application.

Dr. Bob Kuhn moderated a question and answer (Q&A) session with the audience. Question topics included matching the parallel model to the architecture, the types of tools that are available to debug hybrid codes, parallel applications with high I/O requirements, and using dynamic multithreading for MPI processes with different workloads.

"We'd like to tell you that applications do work by invoking these parallel programming models, and the experiences of these people have been that it works easier than you might think."





Dr. Bob Kuhn of Kuck & Associates, Inc., organized and moderated the panel discussion.

Tutorials

Drs. Clay Breshears and Henry Gabb conducted a full-day tutorial entitled "Concurrent Programming with Pthreads." The tutorial covered classic concurrent programming models and problems as well as the 14 core Pthreads functions most useful to scientific programming. Condition variables and thread signaling were also discussed. Examples of Pthreads in scientific computing included multithreaded algorithms from the Command, Control, Communication, and Intelligence Benchmarks as well as matrix multiplication and LU decomposition.

Approximately 35 students attended the tutorial. Many of the questions from the students focused on the differences between Pthreads and OpenMP and the situation where one is more appropriate than the other.

The tutorial notes and quizzes are available on the ERDC MSRC web site at www.wes.hpc.mil under "Announcements."

In the Pthreads tutorial, Dr. Clay Breshears explains monitors—computer science theoretical models for mutual exclusion. SC99 featured more than 20 tutorials with a record number of attendees. Tutorial topics comprised a wide range of topics including data mining, cryptography, and building supercomputers from commercial, off-the-shelf technology.

Reflections of SC99



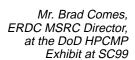


The DoD HPC Modernization Program Exhibit at SC99

Left to right: Dr. Clay Breshears, ERDC MSRC onsite SPP Tools Lead for PET, Dr. Al Geist, Oak Ridge National Laboratories, Dr. Louis Turcotte, Assistant for Technology, ITL, ERDC, and Dr. Henry Gabb, Director of Scientific Computing, ERDC MSRC



Left to right: Drs. Phu Luong, ERDC MSRC onsite EQM Lead for PET, Henry Gabb, and Clay Breshears





SC2000: The Real Millennium Conference

Dr. Louis Turcotte, SC2000 Chair, enticed the audience at SC99 to be cognizant of the technological breakthroughs sure to be seen at SC2000. According to Dr. Turcotte, the SCxy conferences have a "proud tradition of excellence, and SC2000 will continue those expectations."

Several ERDC MSRC team members have already been working diligently to continue the "tradition of excellence." The following is a list of the ERDC MSRC's SC2000 Committee members and their roles on the committee:

Dr. Louis Turcotte SC2000 Chair

John West Co-Chair eSCape 2000 Steve Jones Co-Chair eSCape 2000 Brad Comes Tutorial Committee

Among the highlights of this year's conference will be eSCape 2000, a demonstration of the possibilities of ubiquitous, untethered (i.e., wireless) access to the National Grid. The National Grid is composed of visualization and storage hardware, HPC hardware, and instrumentation (e.g., the Hubble telescope), all linked. The eSCape 2000 committee is looking for "work in progress"—not commercialized products—to showcase novel use of wireless, mobile devices for HPC. Dr. Turcotte invited all of the SC99 attendees to participate.



Dr. Louis Turcotte, SC2000 Chair, with Ms. Donna Shirley, Keynote Speaker at SC99. The SC2000 committee has already completed over a year of planning, and it will be the final conference of the millennium.

Dr. Louis Turcotte is the ERDC

MSRC Government PET Director,
the Assistant for Technology in the
ERDC Information Technology Laboratory,
and served as SC99 Deputy Chair. He is
General Chair for SC2000 in Dallas, TX.

The conference showcase will be the Venture Village, to be seen in the Exhibit Hall. In the Venture Village, a select number of emerging companies, still in the early stages of venture funding and product development, will showcase and discuss the technologies that will affect the future of the HPC community.

SC2000 will be held in Dallas, TX, November 4-10. It is during the week of the national election, so all attendees are encouraged to plan and vote by absentee ballot.

SC2000 Important Dates

APRIL 10, 2000 — Web Site for Submissions Opens

APRIL 28, 2000 — Submission Deadlines: Technical Papers (Extended Abstracts); Gordon Bell Awards Nominations; Tutorials; Panels; Education Program

MAY 28, 2000 — Applications for Education Program Due

JUNE 26, 2000 — Notifications: Conditional Acceptance for Technical Papers; Finalists for the Gordon Bell Awards; Acceptance for Tutorials; Acceptance for Panels; Acceptance for Education Program

JULY 28, 2000 — Submission Deadlines: Research Gems; HPC Games; Birds of a Feather; Research Exhibits; Exhibitor Forum; eSCape 2000

JULY 28, 2000 — Final Version of Technical Papers Due to Confirm Acceptance

JULY 28, 2000 — Student Volunteer Applications
Due

SEPTEMBER 1, 2000 — Notifications: Acceptance for Research Gems, HPC Games, Birds of a Feather, Research Exhibits, Exhibitor Forum, eSCape 2000

SEPTEMBER 1, 2000 — Final Version of Handouts Due for Tutorials

NOVEMBER 4 - 10, 2000 — Conference Dates

NOVEMBER 6 - 9, 2000 — Exhibition Dates

Meet Your Customer

The ERDC MSRC Customer Assistance Center (CAC) has earned a reputation as an outstanding service provider to DoD HPCMP users. The ERDC MSRC has made customer service a priority by making it the core of its Mission Statement.



Michael J. Gough, Lt Col USAF (Ret) CAC Group Lead

Mike is a computer scientist with more than 28 years of experience. In 1971, he earned his B.A. in mathematics from the University of Evansville in Indiana and received his commission as a Lt Col in the United States Air Force. In 1982, he received an M.S. in computer science with an area of concentration in modeling and simulation at the University of Kansas, having been awarded an Air Force Institute of Technology scholarship.

While in the Air Force, he was a member of the Air Launched Cruise Missile, Short Range Attack Missile, B52, B-1B, B2, Tri-Service Attack Missile, and other aircraft test programs. He also served as Branch Chief of the Joint Strategic Targeting Staff at the HQ Strategic Air Command, Deputy to the Commander of the B-1B Flight Test Program at Dyess AFB in Texas, and Division Chief in the

"Special" Test Directorate at HQ Air Force Operational Test and Evaluation Center.

Mike retired from the Air Force in 1992. His first civilian job was as a Senior Analyst for a DoD Joint Test and Evaluation project that measured the effectiveness and assessed the suitability of concealment equipment and techniques against multispectral air attack. From there he was the Lead Avionics Test and Integration Engineer for the Army's RAH-66 Comanche Helicopter Development Test and Evaluation program.

Since August 1996, Mike has served as the Group Lead for the ERDC MSRC CAC.

Mike is married with three children and is still active in sports as a coach, a participant, and a fan.

Lisa Langford CAC Deputy Lead, Accounts Administrator

Lisa has worked at the ERDC since 1988. She began her career as a tape librarian, advanced to computer operations, and was promoted in January 1993 to become the first Accounts Administrator at the ERDC MSRC.

As the Accounts Administrator for the ERDC MSRC, Lisa is responsible for all aspects of account administration, from creation to deletion, and administration of SecurID/PinPad cards. She has played a vital role in the annual account renewal process and the recent implementation of the ERDC MSRC's Accounts, Allocations, and Utilization (AAU) account management system. During times of high user calling, she also answers problem calls.

Lisa provided support to the ERDC MSRC Cray Y-MP users. She had the honor of loading the first 13 users onto the ERDC MSRC Cray C90 and was there to see both the Y-MP and C90 decommissioned.

Lisa is an accomplished equestrian and enjoys both showmanship and running events. She devotes much time to charity work. Lisa spends any remaining spare time with her family, horseback riding, boating, and jet skiing. Lisa is married with one child.





Kelly Lanier Allocation Manager, Customer Assistance Specialist

Kelly began her career at ERDC as part of a student work program in 1989, working full-time and taking night classes toward her B.S. in computer science. Kelly joined the ERDC MSRC family in 1996 as part of the CAC. She provides technical support to users, and 2 years ago, she assumed the additional responsibility of Allocation Manager. She is responsible for updating and tracking the allocation for the users' projects and is a major player during the annual Account Renewal. Her other ancillary activities consist of mailing training notices and creating and maintaining the SecurID cards. Kelly's spare time is consumed by her 7-year-old, and she likes to read and walk.

Assistance Team

Members of the CAC team provide customer support on the myriad issues regarding accounts, applications, and systems. Please remember that your feedback is critical to improving our service to you.



Robin Phillips Customer Assistance Specialist

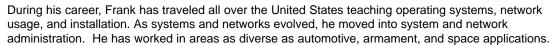
Robin first came to Vicksburg and the ERDC in 1986, serving first as librarian for the 30,000+ volume Magnetic Tape Library, then as computer operator for the ERDC ITL. In 1989, Robin was among the first people selected for the computer operations staff for the newly established Army Supercomputer Center.

Robin has a B.S. in computer science and has served in the CAC since 1993, providing technical support to the ERDC MSRC users. She also serves as the backup Accounts Administrator, knowledgeable in all aspects of account administration, from creation to deletion. She has played a vital role in both the annual account renewal process and the recent implementation and ongoing refinement of the ERDC MSRC's AAU account management system.

In her spare time, Robin is an avid reader and enjoys photography, restoring old furniture, and charitable activities.

James F. Green (Frank) Customer Assistance Specialist

Frank joined the ERDC MSRC CAC in 1998. Although he is one of the newest members of the team, Frank Green is the "gray beard" of the CAC. He started his 30+-year career in computers after receiving a B.S. in electrical engineering from North Carolina State University. His career has spanned the days from "big iron" mainframes to the present. He maintained operating systems in the days before "patches," when programmers actually "diddled" ones and zeros.



As Frank will tell you, going from the early mainframes and proprietary operating systems to the current parallel and "open" operating systems has been an interesting journey. As personal computers began to make their inroads in the industry, Frank taught PC usage to some of the first users. Training is still one of his favorite areas, and he thoroughly enjoys a willing student.

As a CAC team member, Frank handles many of the Kerberos-related problems and file cleanups when users leave or HPC systems are removed. In general, his broad range of expertise allows him to address just about any problem that surfaces in the MSRC.

Frank's spare-time activities include reading, home maintenance, and hockey. The recent establishment of a hockey team in Jackson, MS, has caught his interest.



Melissa Hampton Customer Assistance Specialist

Melissa is the newest member of the CAC team. She began her career at ERDC as part of a student work program in 1986. She developed an extensive Dbase programming project that was used by Corps planners as they studied the recreational use and erosion of the coast of Louisiana. She also worked with collecting and coding water quality data and operating computer models. She later attended Cameron University in Oklahoma and then Mississippi College. During this period, she returned to the ERDC student work program and worked with a Coastal Ecology Group. Her favorite assignment was monitoring endangered sea turtles in the Florida Keys. She also worked with the National Oceanic Atmospheric Administration's Coastal Ocean Program CoastWatch.

Melissa later became a computer operator at ERDC. In May 1999, she became a member of the ERDC MSRC CAC team. In addition to her work with the customers, she is responsible for assembling and assessing the HPC system downtime reports.

Melissa enjoys spending time boating and fishing with her husband and their four children. She also enjoys reading and working with crafts.

Assign GUI

Dr. Mark Fahey, Rebecca Fahey

The CMG at the ERDC MSRC has developed a Graphical User Interface (GUI) for the **assign** command. The **assign** command is commonly used to affect I/O statements allowing programs to read Cray-format binary files (generated on C90s) on SGI Origin 2000 and Cray T3E platforms, which use IEEE binary format. The **assign** command associates Fortran

logical units and file names with I/O processing options like its appropriate binary format.

The GUI-based **assign** command is intended to overcome the otherwise confusing and difficult-to-use nomenclature. The **assign** command has over 20 options that can be used for a variety of purposes. Furthermore, on SGI Origin 2000 platforms, an environment variable FILENV must be set prior to use. In addition, leftover **assign** associations may affect current computing sessions.

With the **assign** GUI, the user sees all **assign** associations on start-up. For SGI IRIX, the GUI shows the current value of the FILENV variable and the location of the .assign file. Furthermore, the **assign** GUI can be used to set new file associations in a point-and-click interface, rather than the confusing command-line interface. The GUI has been developed in Perl/Tk.



A screen view of the prototype **assign** GUI shows the current settings in the **.assign** file. The right side of the window shows the **assign** association for Unit 5 and some of the possible I/O options that could be used. The **assign** associations can be generated or deleted within the GUI.

Answers to quiz on page 12

SGI Origin 2000 (c)

The SGI Origin 2000 is a sharedmemory architecture, but it uses nonuniform memory access rather than a simpler "flat-shared" memory architecture. In a "flat-shared" memory architecture (e.g., SGI PowerChallenge, Sun E10000), all



processors have equal access to all memory locations by way of a memory bus. Parallel programming is easier on flat-shared memory architectures because the programmer does not have to worry about efficient data placement. However, current bus technology limits these systems to a few tens of processors. The SGI Origin 2000 scales to hundreds of processors. This system is best suited to programs that require incremental parallelization (using OpenMP, for example), applications that exhibit task-level concurrency, or applications that are difficult to program efficiently using message passing (e.g., particle dynamics).

The SGI Origin 2000 at the ERDC MSRC has 128-CPUs and 64 GBytes of memory.

IBM SP (b)

The IBM SP uses a distributedmemory architecture of loosely coupled compute nodes. Each node has its own CPU, memory, disk, and operating system. The nodes exchange data through a high-speed switch.



The IBM SP is best suited to coarse-grained parallel applications that require large amounts of memory and are not communication bound. Embarrassingly parallel computations (e.g., Monte Carlo simulations, parametric searches) are suited to the IBM SP.

ERDC MSRC has two IBM SP systems. The larger system has 255 CPUs and 256 GBytes of memory. The smaller system has 127 CPUs and 63 GBytes of memory.

Connecting to an MSRC HPC System via a Palm™ Organizer

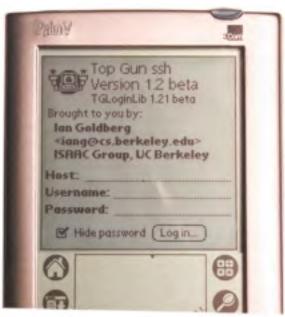
Rebecca Fahey

The use of portable computers with HPC has increased, with a significant number of users taking advantage of the ability to log in to ERDC MSRC HPC systems while away from their office workstations. Laptop computers are frequently used for this purpose, but palmtop computers are gaining popularity. Although the smaller screens of palmtop computers are somewhat limiting, they can be used to access ERDC MSRC HPC accounts, submit or check the status of jobs, issue commands, and perform other simple tasks.

Because palmtop computers are gaining popularity, the ERDC MSRC is taking steps to improve the application of this technology to HPC environments. A guide, "User's Guide: Connecting to ERDC MSRC HPC Systems with a PalmTM Organizer" is available on the ERDC MSRC web site (www.wes.hpc.mil under Publications, Technical Reports, TR 00-06). The guide details the downloading and installation of Top Gun ssh, a free, secure shell program that can be used to log in to an ERDC MSRC HPC system. The configuration of the organizer and an example login session are also included in the guide. Although a PalmTM Organizer is not a replacement for a workstation, or even a laptop, an organizer is useful for short computing sessions away from the normal workstation.

Other efforts to improve mobile computing are under consideration, such as (a) job submission via e-mail along with e-mail receipts containing batch queue information and (b) UNIX command output formatted for the smaller PalmTM screens.

Questions/comments concerning this capability, along with suggestions for improving mobile computing support, may be referred to the ERDC MSRC CAC at 1-800-500-4722 or 1-601-634-4400, Option 1, or by e-mail (*info-hpc @wes.hpc.mil*). Anyone wishing to access ERDC MSRC HPC systems using a palmtop computer are required to contact the ERDC MSRC CAC for access controls.



Users can now access ERDC MSRC HPC via Palm™ organizers. Shown here is the start-up screen of Top Gun ssh when connecting to the ERDC MSRC HPC systems.

Cray T3E (a)

The Cray T3E uses a distributed-memory architecture. Each processing element (i.e., processor and associated memory) is connected to its six nearest neighbors in a 3-D toroidal mesh. This gives the Cray T3E extremely fast communications.



The Cray T3E is best suited to messagepassing applications that are communications bound, have regular communication patterns, and/or make extensive use of collective communication routines (e.g., broadcast, reduce, scatter/gather). Adaptive mesh refinement, for example, is well-suited to the Cray T3E because it is communication intensive and because it involves fine-grain communication between neighboring processors as well as collective communications.

The Cray T3E at the ERDC MSRC has 544 CPUs and 139 GBytes of memory.

IBM Power3 SMP (d)

The IBM Power3 SMP is a distributed, shared-memory hybrid architecture of loosely coupled SMP nodes. Each node has multiple CPUs that share the same memory, disk, and operating system. The SMP nodes exchange data through a high-speed switch.



This architecture is ideally suited to applications with multiple opportunities for parallelism. For example, MPI programs in which the MPI processes can be multithreaded should perform well on the IBM Power3 SMP architecture.

The IBM Power3 SMP at the ERDC MSRC has 64, 8-CPU SMP nodes for a total of 512 CPUs and 256 GBytes of memory.

SuperLU Version 2 Is Now Available

Tom Oppe, Ph.D.

The sparse matrix solver package SuperLU, Version 2, is now available on ERDC MSRC parallel computers. This version of SuperLU comprises three separate packages:

- ➤ Sequential SuperLU is designed for sequential processors with memory hierarchies found in typical cache-based computers.
- ➤ Multithreaded SuperLU (SuperLU_MT) is designed for shared-memory multiprocessors such as the SGI Origin 2000 and the IBM Power3 SMP platforms.
- ➤ Distributed SuperLU (SuperLU_DIST) is designed for distributed-memory parallel computers such as the Cray T3E and the IBM SP2 platforms. SuperLU_DIST is new with Version 2.

SuperLU is a package designed to solve large sparse sets of linear equations using variants of the Gaussian elimination algorithm. The matrices corresponding to the linear system may be either symmetric or nonsymmetric in coefficients and in nonzero structure. Various algorithms are used to reorder the matrix to reduce the fill-in of nonzeros during factorization and, in the case of SuperLU_MT and SuperLU_DIST, increase the parallel performance of the factorization step. In addition, a dynamic threshold-pivoting algorithm is used to improve the numerical stability of the factorization for SuperLU and SuperLU_MT. SuperLU_DIST uses a static reordering algorithm (i.e., done prior to factorization) to increase numerical stability.

The choice of orderings includes (a) the natural ordering (i.e., as supplied by the user), (b) Multiple Minimum Degree applied to the structure of A^T * A,

(c) Multiple Minimum Degree applied to the structure of $A + A^T$, and (d) Column Approximate Minimum Degree (COLAMD). The COLAMD ordering is new to Version 2 and is often faster than the older orderings.

In addition to reordering algorithms, SuperLU contains routines to equilibrate (i.e., diagonally scale) the matrix to improve and estimate its condition number, factor the matrix, solve for one or more right-hand sides, iteratively refine the solution, and compute error bounds for the solution.

Sequential SuperLU is available in all four precisions: single- and double-precision real-valued and single- and double-precision complex-valued (denoted S, D, C, and Z, respectively). SuperLU_MT is available only in the D precision, and SuperLU_DIST is only available in the D and Z precisions. The source code is written in the C programming language, with the matrices entered as C programming structures; therefore, a Fortran-to-C routine must be used when calling SuperLU from Fortran. The bridge routine requires that the matrix be stored in the Harwell-Boeing sparse matrix format. If you need help, call the ERDC CAC, and we will be glad to work with you on using SuperLU.

The SGI Origin 2000 SuperLU libraries are in /usr/local/usp/PETtools/lib64 and /usr/local/usp/PETtools/lib32 for use with programs compiled with -64 and -n32, respectively. The names of the libraries are as follows:

libSuperLU.a libSuperLU_MT.a libSuperLU_DIST.a

for sequential SuperLU for multithreaded SuperLU_MT for distributed SuperLU_DIST

More information

www.nerse.gov/~xiaoye/SuperLU/
SuperLU source codes and Users' Guide
www.nerse.gov/~xiaoye/.

Several papers describing the algorithms used in SuperLU

ERDC MSRC Computer Systems Are Gems

Rose Dykes

Why have names for our equipment? Just as our homes have addresses, the digital world uses Internet Protocol addresses (long numbers) to route information among computer systems. Who wants to be bothered with trying to remember numbers? Naming our computers solves this problem

In the early days of the MSRC, system names did not follow any consistent logic. Dennis Gilman, the ERDC MSRC's Government Contracting Officer's Representative, is now responsible for naming the computer systems, storage facilities, and workstations at the ERDC MSRC. The first step was to establish the criteria for naming systems. Dennis' approach was

to select a theme that supported names that are both hierarchical and globally recognized. Names that are easily pronounced and spelled are selected in conjunction with the theme.

The names chosen for the large computer systems follow the theme of precious gems such as Cobalt, Diamond, Emerald, and Ruby. Precious metals, for example, Gold, Silver, and Platinum, are used to name storage devices. Lesser metal names such as Copper, Bronze, and Brass are used for workstations. We even have a small desktop system named Dirt. Look for more gems in the future of the ERDC MSRC.



Dennis Gilman with Cobalt, the IBM Power3 SMP

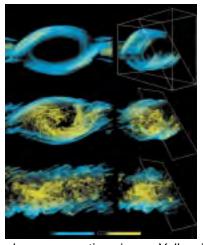
interview with . . . (continued from page 7)

the highest resolution and largest simulations of stratified wind shear that have ever been conducted. The solutions are a significant contribution to the ABL project. They can also be used to develop better SGS techniques for stratified turbulence so that method (a) might be reliably used in the future to simulate atmospheric turbulence.

How large have your problems grown over the past **5 years?** We are now in our third year of com-puting at the DoD centers within the High Performance Computing Modernization Office Challenge program. Before that, the project was supported by the NSF Grand Challenge program, primarily at the Pittsburgh Supercomputing Center, but also in San Diego and Illinois. Five years ago, as a part of the NSF effort, my largest simulations used 384 x 384 x 128 spectral modes running on a Cray C90. They required 16 processors and consumed less than 1,000 CPU hours. With the resources now available through the DoD Challenge program, my largest simulations use 1,000 x 350 x 2,000 spectral modes and require 500 Cray T3E processors. Between 100,000 and 200,000 CPU hours are needed to conduct such a simulation, and several terabytes of numerical data are generated. This represents a more than 100-fold increase in compute capability from 5 years ago. Needless to say, the ability to substantively contribute to a deeper scientific understanding of turbulence dynamics and effects are significantly advanced.

Do you anticipate them growing? There are two major sources for turbulence in the stratosphere: (a) wind shear and (b) overturning gravity waves. For the past 2 years, we have focused on wind shear. In the coming years, we will increasingly turn our attention to gravity-wave breaking. We expect that the gravity-wave breaking solutions will require physical domains several times the size of the wind-shear solutions we are currently computing. Also, through our research and that of others, it appears that alternate paths to turbulence are possible for wind shear. It turns out that of the two dominant routes, the case we have examined so far is associated with the shorter length scale. As a result, future simulations investigating the other route to turbulence in wind shear necessarily require longer domains. Therefore, in order to further our understanding of the various important paths turbulence can take in the atmosphere, we do indeed anticipate that our problem sizes will grow.

What is the impact on the DoD? Aside from the direct impact our work has on the ABL program, the solutions we have computed so far within the DoD Challenge program possess potentially far-reaching



Evolution of turbulence in a simulation of wind shear conducted on as many as 500 processors using the Cray T3E at the ERDC MSRC. Panels on the left show side views of developing wind shear. Right panels

show perspective views. Yellow indicates high turbulence intensity, i.e., strong spatial and temporal fluctuations of air velocity. Blue shows large variations in the optical properties of the flow, e.g., large variations in the index of refraction gradient.

applications. For example, because (a) we have neglected no terms in the incompressible fluid equations and (b) we employ extremely accurate spectral techniques for their solution, the numerical solutions represent our best and most complete description of turbulence in stratified fluids. As such, they could and should be used to develop and test improved SGS models for turbulence in stratified fluids. Such a development would obviously benefit the ABL program because even larger scale simulations would be possible. Furthermore, such an effort would also pay off for both ocean modeling and weather prediction, as both the ocean and the atmosphere possess significant stratification effects that are currently poorly modeled with existing SGS schemes.

Recent publications

Werne, J., and D. C. Fritts. 1998. "Turbulence in stratified and sheared fluids: T3E simulations." DoD HPCMO Users Symposium.

Werne, J., and D. C. Fritts. 1999. "Stratified shear turbulence: Evidence of anisotropy." DoD HPCMO Users Symposium.

Werne, J., and D. C. Fritts. 1999. "Stratified shear turbulence: Evolution and statistics," *Geophysical Research Letters*, Vol 26, 439-442 (cover).

Gibson-Wilde, D., J. Werne, D. C. Fritts, and R. J. Hill. 1999. "Direct numerical simulation of VHF radar measurements of turbulence in the mesosphere," *Radio Science* (in press).

Hill, R. J., D. Gibson-Wilde, J. Werne, and D. C. Fritts. 1999. "Turbulence-induced fluctuations in ionization and application to PMSE," *Earth Planets Space*, Vol 51, 499-513.

Werne, J., and D. C. Fritts. 1999. "Turbulence and mixing in a stratified shear layer: 3D K-H simulations at Re=24,000," *Physics and Chemistry of the Earth* (in press).

Fritts, D. C., J. Werne, and T. L. Palmer. 2000. "Dynamics of turbulence in stratified shear flows at high Reynolds number" (submitted for publication).

Below is a list of acronyms commonly used among the DoD HPC community. You will find these acronyms throughout the articles in this newsletter.

AAU Administration, Allocation, and Utilization

ABL AirBorne Laser
ALU Arithmetic-Logic Unit
ASC Aeronautical Systems Center
CAC Customer Assistance Center

CD Compact Disk

CHSSI Common High Performance Computing Software Support Initiative

CMG Computational Migration Group

CPU Central Processing Unit

CSC Computer Sciences Corporation

DoD Department of Defense

EQM Environmental Quality Modeling and Simulation ERDC Engineer Research and Development Center

GL Geotechnical Laboratory
GUI Graphical User Interface
HPC High-Performance Computing

HPCMP High Performance Computing Modernization Program

InfoVis Information Visualization

I/O Input/Output

IT Information Technology

ITL Information Technology Laboratory

JPL Jet Propulsion Laboratory
MPI Message-Passing Interface
MSRC Major Shared Resource Center
MSU Mississippi State University
NAVO Naval Oceanographic Office
NSF National Science Foundation
PBS Portable Batch System

PET Programming Environment and Training
PVG Parallel Visualization and Graphics

Q&A Questions and Answers SC Supercomputing

SGS Subgrid Scale

SMP Symmetric Multiprocessing
SPP Scalable Parallel Programming
WPAFB Wright-Patterson Air Force Base

The contents of this newsletter are not to be used for advertising, publication, or promotional purposes. Citation of trade names does not constitute an official endorsement or approval or the use of such commercial products.

The ERDC MSRC welcomes comments and suggestions regarding *The Resource*. Please send submissions to the following e-mail address:

info-hpc@wes.hpc.mil

In Perspective

- Seventeen and one-half miles of wiring was required to connect all of the components associated with the ERDC MSRC's Uninterrupted Power Supply (UPS).
- ➤ The UPS facility consists of three Caterpillar engines. The slab of concrete supporting one Caterpillar engine is 8 ft thick.
- Two hundred gallons of oil is needed to fill one Caterpillar engine, and 115 gal of antifreeze is required to fill one radiator.

- The UPS consumes 420 gal of fuel per hour.
- ➤ The Origin 2000 weighs 6,300 lb; the Cray T3E weighs 5,600 lb; and the IBM complex (both SPs and the Power3 SMP) weighs 31,200 lb.
- ➤ The ERDC MSRC HPC systems can theoretically perform up to 1.3 trillion calculations per second.



Technical Reports

99-35	"Programming Environment and	Training Annual Report - Year Three"
-------	------------------------------	--------------------------------------

00-01 Richard Weed, "Building Multidisciplinary Applications with MPI"

00-02 Richard Weed, "MDARUN A Package of Software for Creating Multidisciplinary Applications with MPI"

00-03 Mark R. Fahey, "Locating Floating-Point Exceptions on the SGI Origin 2000"

00-04 David Littlefield and J. Tinsley Oden, "Implementation of Adaptive Mesh Refinement into an Eulerian Hydrocode"

These technical reports can be accessed at http://www.wes.hpc.mil/.



ERDC MSRC Newsletter Editoral Staff

Interdisciplinary Coordinator

Mary L. Hampton

Chief Editors

Rose Dykes Mary Gabb

Visual Information Specialist

Betty Watson

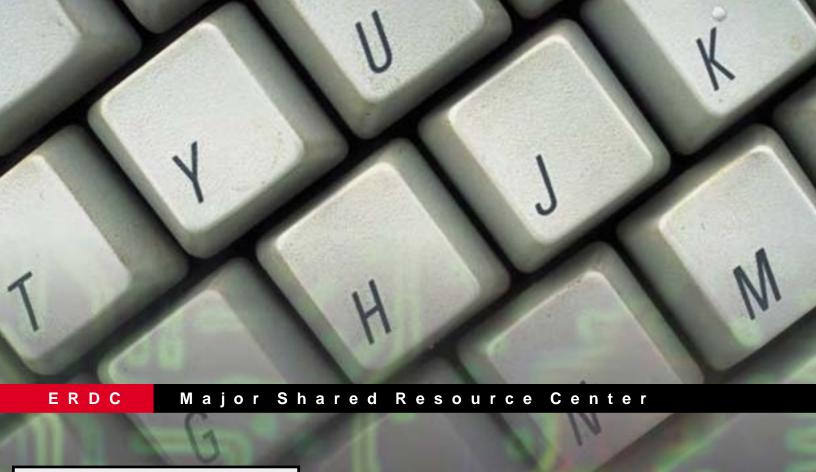
ERDC MSRC Customer Assistance Center Web site: http://www.wes.hpc.mil/ E-mail: info-hpc@wes.hpc.mil Telephone: 1-800-500-4722

Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the DoD. Your comments, ideas, and contributions are welcome.

Design and layout provided by the Visual Production Center, Information Technology Laboratory,

U.S. Army Engineer Research and Development Center.

Approved for public release; distribution is unlimited.



U.S. Army Engineer Research and Development Center ATTN: CEERD-IH (Rose Dykes)
3909 Halls Ferry Road
Vicksburg, MS 39180-6199